

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	7142	(707/8,200,201,202,203,204).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/08/18 18:46
L2	1407	(711/100).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/08/18 18:47
L3	134	(714/100).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/08/18 18:48
L4	8611	(1 or 2 or 3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 18:48
L5	373	"consistent copy"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 18:48
L6	274	"current snapshot"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 18:49
L7	3	5 and 6	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:01
L8	4991	"database object"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:01

EAST Search History

L9	734	"database object".clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:01
L10	627	snapshot near3 updat\$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:02
L11	83	snapshot near3 updat\$3.clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:02
L12	1	remov\$3 near5 11	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:04
L13	334	"log entry".clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:05
L14	5485	"start time".clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:05
L15	8	13 and 14	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:51
L16	1803	(711/161,162).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/08/18 19:53

EAST Search History

L17	80	10 and 16	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:55
L18	3	17 and 5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/08/18 19:55



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **snapshot update**Found **5** of **184,245**

Sort results by

Display results


[Save results to a Binder](#)

[Search Tips](#)
☐ Open results in a new window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 5 of 5

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [On network-aware clustering of Web clients](#)



Balachander Krishnamurthy, Jia Wang

 August 2000 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication SIGCOMM '00**, Volume 30 Issue 4

Publisher: ACM Press

Full text available: pdf(568.99 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Being able to identify the groups of clients that are responsible for a significant portion of a Web site's requests can be helpful to both the Web site and the clients. In a Web application, it is beneficial to move content closer to groups of clients that are responsible for large subsets of requests to an origin server. We introduce clusters---a grouping of clients that are close together topologically and likely to be under common administrative control. We identify clu ...

2 [A snapshot differential refresh algorithm](#)



Bruce Lindsay, Laura Haas, C. Mohan, Hamid Pirahesh, Paul Wilms

 June 1986 **ACM SIGMOD Record , Proceedings of the 1986 ACM SIGMOD international conference on Management of data SIGMOD '86**, Volume 15 Issue 2

Publisher: ACM Press

Full text available: pdf(827.36 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This article presents an algorithm to refresh the contents of database snapshots. A database snapshot is a read-only table whose contents are extracted from other tables in the database. The snapshot contents can be periodically refreshed to reflect the current state of the database. Snapshots are useful in many applications as a cost effective substitute for replicated data in a distributed database system. When the snapshot contents are a simpl ...

3 [Measuring the internet's vital statistics: Collecting the internet AS-level topology](#)



Beichuan Zhang, Raymond Liu, Daniel Massey, Lixia Zhang

 January 2005 **ACM SIGCOMM Computer Communication Review**, Volume 35 Issue 1

Publisher: ACM Press

Full text available: pdf(307.85 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

At the inter-domain level, the Internet topology can be represented by a graph with Autonomous Systems (ASes) as nodes and AS peerings as links. This AS-level topology graph has been widely used in a variety of research efforts. Conventionally this topology graph is derived from routing tables collected by Route Views or RIPE RIS. In this work, we assemble the most complete AS-level topology by extending the conventional method along two dimensions. First, in addition to using data from RouteVie ...

4 The design of POSTGRES



Michael Stonebraker, Lawrence A. Rowe

June 1986 **ACM SIGMOD Record , Proceedings of the 1986 ACM SIGMOD international conference on Management of data SIGMOD '86**, Volume 15 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.91 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents the preliminary design of a new database management system, called POSTGRES, that is the successor to the INGRES relational database system. The main design goals of the new system are to provide better support for complex objects, provide user extendibility for data types, operators and access methods, provide facilities for active databases (i.e., alerters and triggers) and inferencing including forward- ...

5 TeamRooms: network places for collaboration



Mark Roseman, Saul Greenberg

November 1996 **Proceedings of the 1996 ACM conference on Computer supported cooperative work**

Publisher: ACM Press

Full text available: [pdf\(1.07 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: GroupKit, groupware, shared electronic spaces

Results 1 - 5 of 5

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)

Full text available:  pdf(441.35 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Failure-resilient, scalable, and secure read-write access to shared information by mobile and static users over wireless and wired networks is a fundamental computing challenge. In this article, we describe how the Coda file system has evolved to meet this challenge through the development of mechanisms for server replication, disconnected operation, adaptive use of weak connectivity, isolation-only transactions, translucent caching, and opportunistic exploitation of hardware surrogates. For each ...

Keywords: Adaptation, Linux, UNIX, Windows, caching, conflict resolution, continuous data access, data staging, disaster recovery, disconnected operation, failure, high availability, hoarding, intermittent networks, isolation-only transactions, low-bandwidth networks, mobile computing, optimistic replica control, server replication, translucent cache management, weakly connected operation


4 [An on-the-fly reference counting garbage collector for Java](#)



Yossi Levroni, Erez Petrank

October 2001 **ACM SIGPLAN Notices , Proceedings of the 16th ACM SIGPLAN conference on Object oriented programming, systems, languages, and applications OOPSLA '01**, Volume 36 Issue 11

Publisher: ACM Press

Full text available:  pdf(280.30 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Reference counting is not naturally suitable for running on multiprocessors. The update of pointers and reference counts requires atomic and synchronized operations. We present a novel reference counting algorithm suitable for a multiprocessor that does not require any synchronized operation in its write barrier (not even a compare-and-swap type of synchronization). The algorithm is efficient and may complete with any tracing algorithm.

5 [Correctness-preserving derivation of concurrent garbage collection algorithms](#)



Martin T. Vechev, Eran Yahav, David F. Bacon

June 2006 **ACM SIGPLAN Notices , Proceedings of the 2006 ACM SIGPLAN conference on Programming language design and implementation PLDI '06**, Volume 41 Issue 6

Publisher: ACM Press

Full text available:  pdf(251.26 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Constructing correct concurrent garbage collection algorithms is notoriously hard. Numerous such algorithms have been proposed, implemented, and deployed - and yet the relationship among them in terms of speed and precision is poorly understood, and the validation of one algorithm does not carry over to others. As programs with low latency requirements written in garbagecollected languages become part of society's mission-critical infrastructure, it is imperative that we raise the level of confidence ...

Keywords: concurrent algorithms, concurrent garbage collection, synthesis, verification


6 [Reducing pause time of conservative collectors](#)



Toshio Endo, Kenjiro Taura

June 2002 **ACM SIGPLAN Notices , Proceedings of the 3rd international symposium on Memory management ISMM '02**, Volume 38 Issue 2 supplement

Publisher: ACM Press


Full text available:  pdf(182.62 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes an incremental conservative garbage collector that significantly reduces pause time of an existing collector by Boehm et al. Like their collector, it is a true conservative collector that does not require compiler cooperation but uses virtual memory primitives (page protection) of operating systems for write barriers. While much

successful work has been done on incremental collectors in general, achieving small pause time of the order of a few milliseconds in such uncooperat ...

Keywords: concurrent garbage collection, conservative garbage collection, memory management, parallel garbage collection

7 User evaluation of Físchlár-News: An automatic broadcast news delivery system

 Hyowon Lee, Alan F. Smeaton, Noel E. O'connor, Barry Smyth

April 2006 **ACM Transactions on Information Systems (TOIS)**, Volume 24 Issue 2


Publisher: ACM Press

Full text available:  pdf(1.25 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Technological developments in content-based analysis of digital video information are undergoing much progress, with ideas for fully automatic systems now being proposed and demonstrated. Yet because we do not yet have robust operational video retrieval systems that can be deployed and used, the usual HCI practise of conducting a usage study and an informed iterative system design is thus not possible. Físchlár-News is one of the first automatic, content-based broadcast news analys ...

Keywords: User-evaluation, content-based video retrieval, usage analysis

8 New garbage collection algorithms and strategies: Garbage-first garbage collection

 David Detlefs, Christine Flood, Steve Heller, Tony Printezis

October 2004 **Proceedings of the 4th international symposium on Memory management**

Publisher: ACM Press

Full text available:  pdf(199.59 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

<i>Garbage-First</i> is a server-style garbage collector, targeted for multi-processors with large memories, that meets a soft real-time goal with high probability, while achieving high throughput. Whole-heap operations, such as global marking, are performed concurrently with mutation, to prevent interruptions proportional to heap or live-data size. Concurrent marking both provides collection "completeness" and identifies regions ripe for reclamation via compacting evacuation. This ev ...


Keywords: concurrent garbage collection, garbage collection, garbage-first garbage collection, parallel garbage collection, soft real-time garbage collection

9 Tool-based approach to distributed database design: includes Web-based forms design for access to academic affairs data

 David A. Owens, Frederick T. Sheldon

February 1999 **Proceedings of the 1999 ACM symposium on Applied computing**

Publisher: ACM Press

Full text available:  pdf(806.36 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: ORACLE, Web, academic affair, database, distributed

10 On bounding time and space for multiprocessor garbage collection

 Guy E. Blelloch, Perry Cheng

May 1999 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1999 conference on Programming language design and implementation PLDI '99**, Volume 34 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.85 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index](#)

[terms](#)

This paper presents the first multiprocessor garbage collection algorithm with provable bounds on time and space. The algorithm is a real-time shared-memory copying collector. We prove that the algorithm requires at most $2(R(1 + 2/k) + N + 5PD)$ memory locations, where P is the number of processors, R is the maximum reachable space during a computation (number of locations accessible from the root set), N is the maximum number of reachable objects, ...

11 [Reality mining: sensing complex social systems](#)

Nathan Eagle, Alex (Sandy) Pentland

March 2006 **Personal and Ubiquitous Computing**, Volume 10 Issue 4

Publisher: Springer-Verlag

Full text available: [pdf\(846.82 KB\)](#) Additional Information: [full citation](#), [abstract](#)

We introduce a system for sensing complex social systems with data collected from 100 mobile phones over the course of 9 months. We demonstrate the ability to use standard Bluetooth-enabled mobile telephones to measure information access and use in different contexts, recognize social patterns in daily user activity, infer relationships, identify socially significant locations, and model organizational rhythms.

Keywords: Bluetooth, Complex social systems, Mobile phones, User modeling, Wearable computing

12 [Characterizing the synchronization behavior of parallel programs](#)

Helen Davis, John Hennessy

January 1988 **ACM SIGPLAN Notices , Proceedings of the ACM/SIGPLAN conference on Parallel programming: experience with applications, languages and systems PPEALS '88**, Volume 23 Issue 9

Publisher: ACM Press

Full text available: [pdf\(1.53 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Contention for synchronization locks and delays waiting for synchronization events can substantially increase the running time of a parallel program. This makes it important to characterize the synchronization behavior of programs and to provide analysis tools to aid both the hardware and software designer in evaluating design alternatives. This paper describes a tracing facility that is incorporated into a synchronization package. This facility provides a portable means to accurately and e ...

13 [Piazza: a desktop environment supporting impromptu and planned interactions](#)

Ellen A. Isaacs, John C. Tang, Trevor Morris

November 1996 **Proceedings of the 1996 ACM conference on Computer supported cooperative work**

Publisher: ACM Press

Full text available: [pdf\(1.67 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: awareness, enterprise-wide communication, informal communication, networkers, unintended interactions

14 [Cases from the field: Where am I and who am I?: issues in collaborative technical help](#)

Michael Twidale, Karen Ruhleder

November 2004 **Proceedings of the 2004 ACM conference on Computer supported cooperative work**

Publisher: ACM Press

Full text available: [pdf\(296.47 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In a study of collaborative help-giving within several organizations settings, we identified two forms of trouble and bewilderment that we explore further in this paper. In one case, the user is confused about where they, their files, or other resources are within a larger technical infrastructure (Where am I?). In the second case, the user isn't sure which login is needed and which actions are allowed (Who am I?). We believe that these issues carry important implications for the design of in ...

Keywords: CSCW, collaborative help-giving, informal learning

Results 1 - 14 of 14

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)